NATIONAL WEATHER SERVICE

Rapid Refresh Forecast System

A case study for using cloud resources to reduce R2O transition times

October 22, 2020

EMC : Arun Chawla, Jacob Carley, James Abeles, Raj Panda
GSL : Curtis Alexander, Christina Holt, Christopher Harrop, Daniel Abdi
NSSL : Louis Wicker

Outline

- NPS Production Suite and UFS
- RRFS
- Using the Cloud for development

Shared Community Infrastructure Support for UFS Development

Infrastructure for data assimilation:
Joint Effort for Data assimilation Integration (**JEDI**)

Infrastructure for coupling models together:

- NOAA Environmental Modeling System (**NEMS**) coupler
- based on the Earth System Modeling Framework (**ESMF**)
- using National Unified Operational Prediction Capability (**NUOPC**) conventions

Infrastructure for interoperable physics:

- Common Community Physics Package (**CCPP**) framework

Infrastructure for Code Management:

- Git based repositories with Gitflow

1. Coupling components
New ESMF/NUOPC mediator (CMEPS/NEMS)

2. Interoperable atmospheric physics
CCPP & CPF frameworks

3. Community-friendly workflow
CIME - CROW unification, CIME Case Control System

4. Hierarchical model development capabilities
Extensions of CIME data models, unit, & system testing

5. Forecast Verification: Comparison to Observations
Extension of MET+

6. Software Repository Management
NCAR manage_externals tool

7. User / Developer Support
DTC and CESM Capabilities

NOAA-NCAR MoA Work Areas

NPS Transitioning to UFS Applications

UFS
UNIFIED FORECAST SYSTEM

FOCUS	local	regional	global		
PREDICTIVE TIME SCALE	hour	day	week	month	year

UFS APPLICATIONS

Application	UFS Application
weather and seasonal	Warn on Forecast System; Rapid Refresh Forecast Systems; Global Forecast System; Sub-Seasonal Forecast System; Seasonal Forecast System
hurricane	Hurricane Analysis and Forecast System
space weather	Whole Atmosphere Model
marine and coastal	Coastal Modeling Systems
air quality	Air Quality Systems
flood and hydrological	National Water Model

Toward a simpler, UFS-based Production Suite (notional)

NPS Modeling System	Current Version	Q3 FY 20	Q4 FY 20	Q1 FY 21	Q2 FY 21	Q3FY 21 - Q2FY22 MORATORIUM	Q3 FY 22	Q4 FY 22	Q1 FY 23	Q2 FY 23	Q3 FY 23	Q4 FY 23	Q1 FY 24	Q2 FY 24	Q3 FY 24	Q4 FY 24	UFS Application
Global Weather & Global Analysis	GFS/ GDASv15				GFSv16									GFSv17/ GEFSv13			UFS Medium Range & Sub-Seasonal
Global Waves	GWMv3																
Global Weather Ensembles	GEFSv11		GEFSv12														
Global Wave Ensembles	GWESv3																
Global Aerosols	NGAC v2																
Short-Range Regional Ensembles	SREFv7																
Global Ocean & Sea-Ice	RTOFSv1.2			RTOFSv2					RTOFSv3								UFS Marine & Cryosphere
Global Ocean Analysis	GODASv2								GODASv3								
Seasonal Climate	CDAS/ CFSv2															SFSv1	UFS Seasonal
Regional Hurricane 1	HWRFv12		HWRFv13				HAFSv1				HAFSv2				HAFSv3		UFS Hurricane
Regional Hurricane 2	HMONv2	HMONv3															
Regional High Resolution CAM 1	HiRes Window v7			HIRESWv8						RRFSv1				RRFSv2			UFS Short-Range Regional HiRes CAM & Regional Air Quality
Regional High Resolution CAM 2	NAM nests/ Fire Wxv4																
Regional High Resolution CAM 3	RAPv4/ HRRRv3			RAPv5/ HRRRv4													
Regional HiRes CAM Ensemble	HREFv2			HREFv3													
Regional Mesoscale Weather	NAMv4																
Regional Air Quality	CMAQv5						CMAQv6										
Regional Surface Weather Analysis	RTMA/ URMA v2.7		RTMA/ URMA v2.8								3DRTMA /URMAv 3						
Atmospheric Transport & Dispersion	HySPLITv7						HySPLIT v8								HySPLIT v9		UFS Air Quality & Dispersion
Coastal & Regional Waves	NWPSv1.2			NWPS v1.3			NWPS v1.4						RWPSv1				UFS Coastal
Great Lakes	GLWUv3.4						GLWUv4								GLWUv5		UFS Lakes
Regional Hydrology	NWMv2			NWMv3							NWMv4						UFS Hydrology
Space Weather 1	WAM/IPEv1															WAMv2	UFS Space Weather
Space Weather 2	ENLILv1																

UFS: Simplification of Regional Model Production Suite

Long Term

Global (10-15 km, 4-24/day)

GFS → GEFS

GFS/GEFS to subsume RAP/NAM/SREF

CONUS/OCONUS CAM (3 km, 2-24/day)

Rapid Refresh Forecast System

RRFS

RRFS to subsume HRRR/NAMnest /HREF

Sub-CONUS CRM (1 km, on-demand)

Warn on Forecast System

WoFS

What is the RRFS and how do we get there?

*Material courtesy of C. Alexander

The Rapid Refresh Forecast System

- Rapid Refresh Forecast System (RRFS)
  - Based on the UFS-Limited Area Model (LAM)
    - Underpinned by the Finite-Volume Cubed-Sphere dynamical core (FV3)
  - Rapidly updated (hourly)
  - Convection-allowing (~3 km)
  - Ensemble data assimilation (est. 30 or 40 members)
  - Ensemble forecasts (est. 10 members)
    - 18h+ hourly
    - 36-60h every 6 to 12 hours
- When? ~FY23
- Exact configuration for data assimilation still under development
- Facilitate replacement of several current regional systems
  - e.g., NAM+nests, RAP +HRRR, HiRes Windows, HREF
- How do we get there? And with what resources?

LAM vs Nest

- FV3 was originally a global model with a nesting capability
  - No limited area component
- Practical reasons drive need for LAM
  - More efficient
  - Focused development
  - Rapidly updating data assimilation, etc.
  - Little impact on forecast quality for short-range forecasts

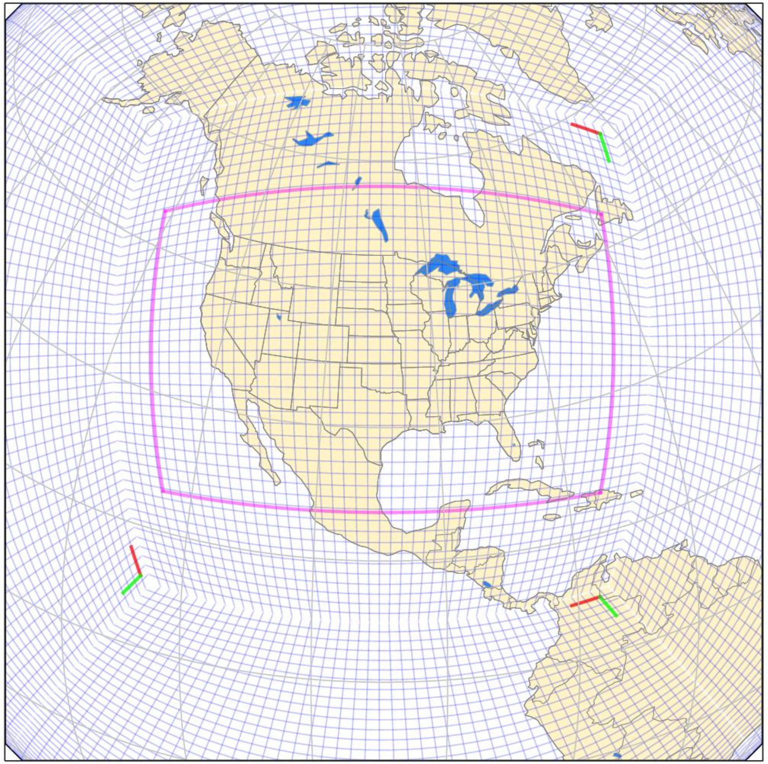

LAM vs Nest Clock Times

- Using NOAA Hera R&D HPC System
- **LAM needs less than half tasks required to get output in *same* amount time as NEST**
- **For a given set of tasks, the LAM completes the forecast in about half the amount of time as NEST**

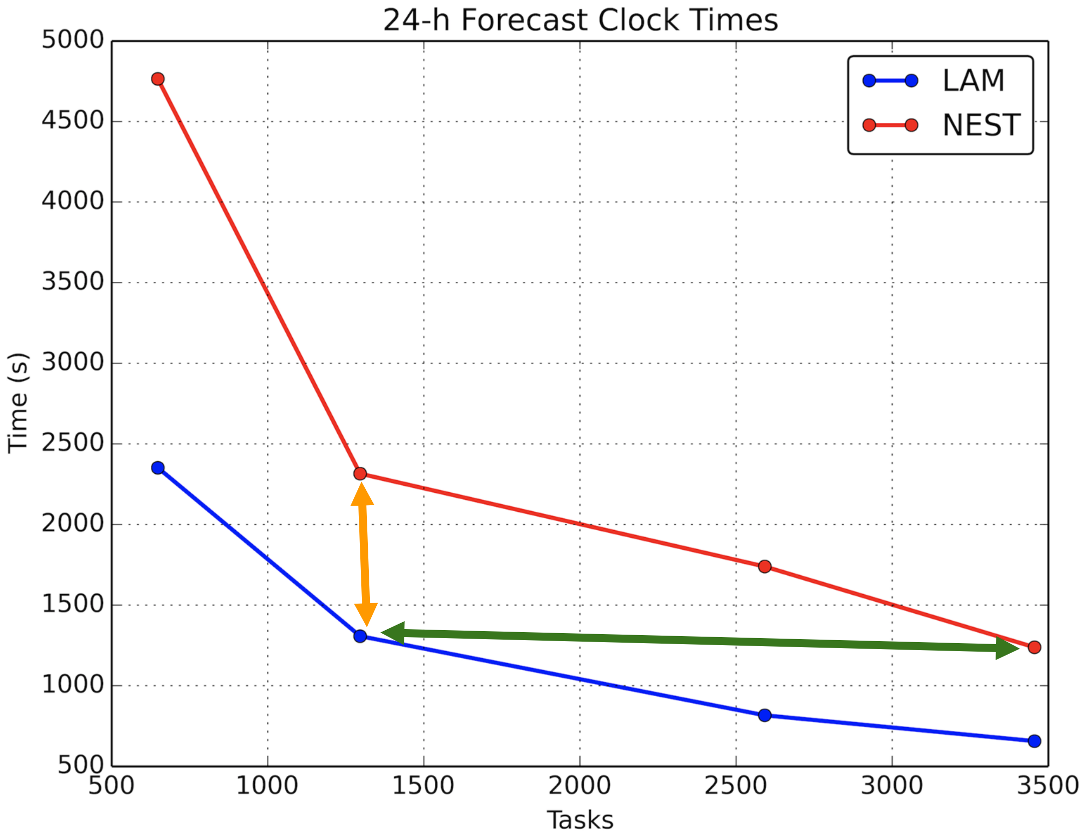

Nest vs LAM for Precip

- Compare Nest and LAM
  - 60 hour forecasts
- Performance Diagram for 24-h accumulated precip (March 15 - April 16 2019)
  - Similar performance across all thresholds

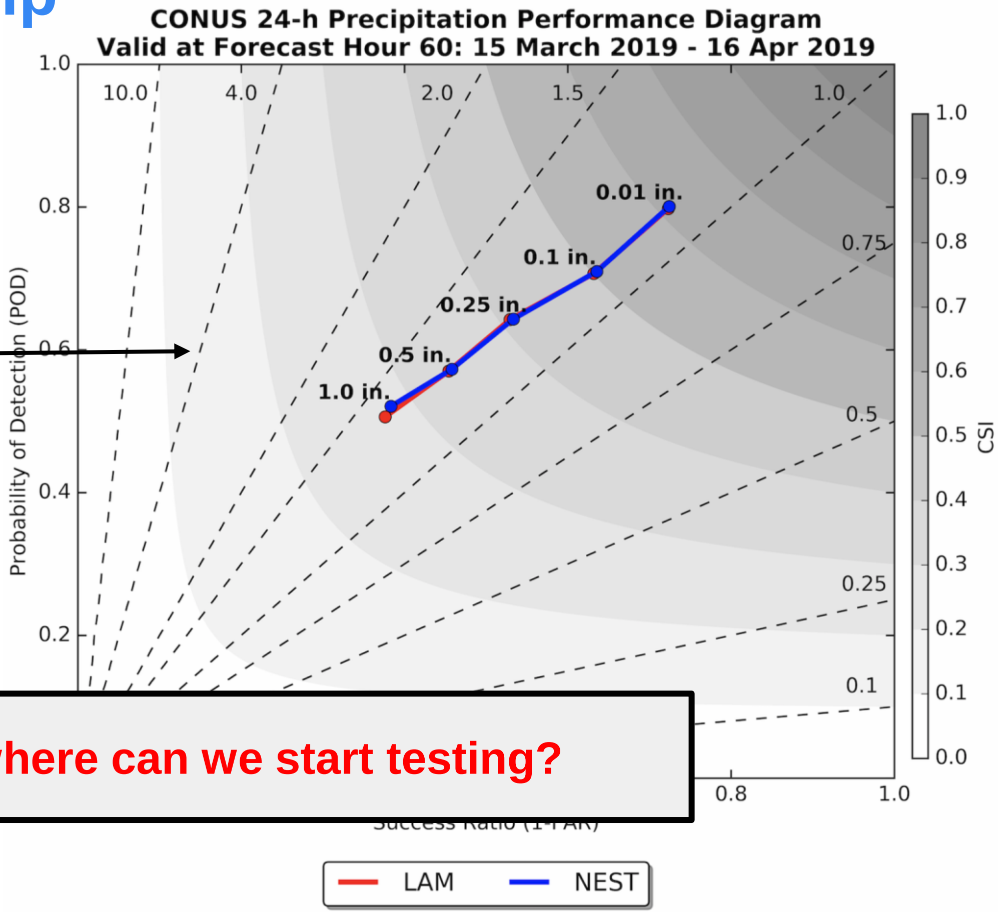

LAM seems to work - so where can we start testing?

RRFS Development Resource Needs

FV3-CAM [RRFS] development *<u>needs</u>*

High water mark on WCOSS Phase 3 Development system

Current usage [*non-CAM*]

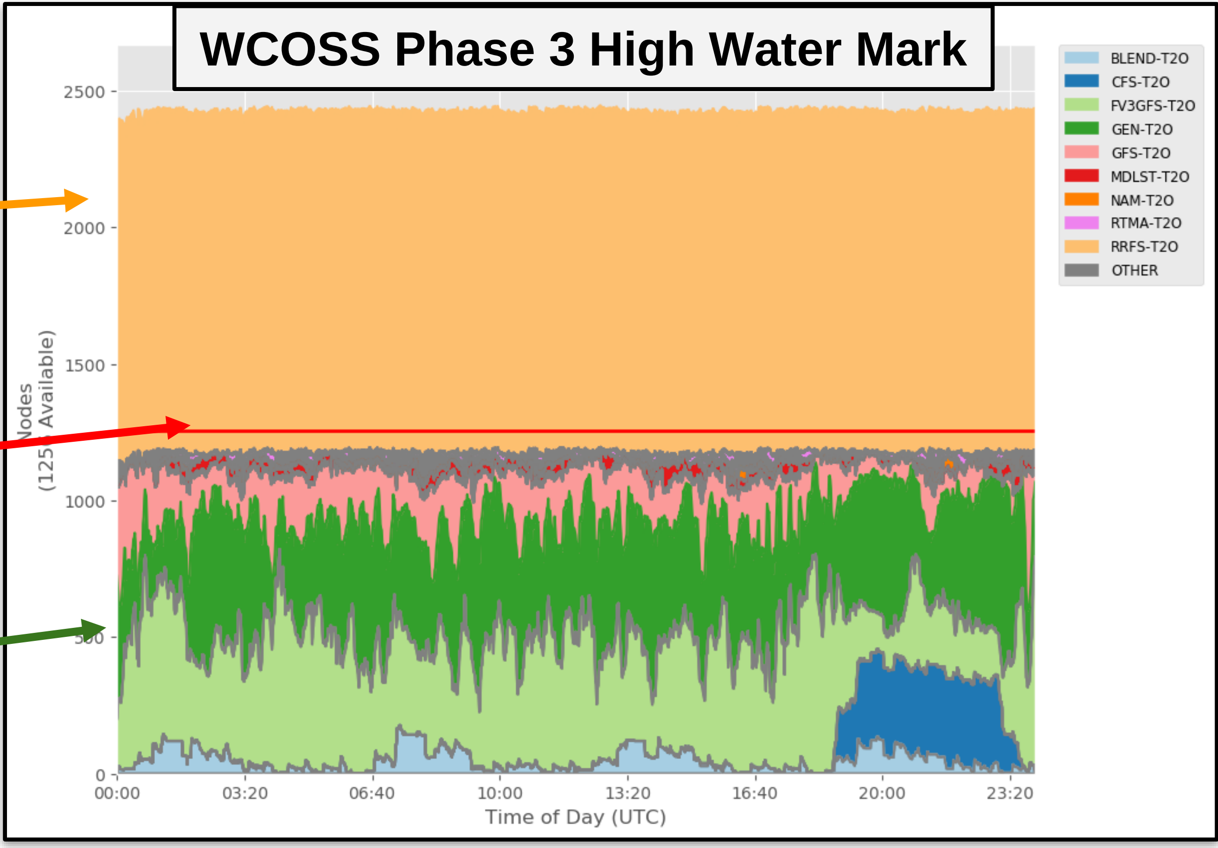

RRFS development requires considerable computational resources that are not currently available to developers

Can We Do this on the Cloud?

- Hurricane Disaster Supplemental Supported Project
- Port, deploy, and test RRFS prototype(s) in the cloud
- Flexible
  - If this works it opens many possible doors for future R&D resource needs, especially for early development that may not need real time testing
- UFS has been run in the cloud before
  - Efforts to put RRFS on cloud will enhance system portability for the UFS infrastructure
- Is cloud HPC slow?
  - **#136** on **June 2019 top500** list is a cloud instance!

			Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)
136	Descartes Labs United States	**Amazon EC2 C5 Instance cluster us-east-1a** - Amazon EC2 Instance Cluster C5, Xeon Platinum 8124M 18C 3GHz, 25G Ethernet Amazon Web Services	41,472	1,926.4	3,981.3

Moving to the Cloud - overview

- Multiple computing components
  - Preprocessing, Data Assimilation, Forecast, Post Processing, Visualization, etc.
- Performance characteristics & requirements
  - Some components (e.g. forecast) will need 10s of nodes with fast interconnect
  - Some (e.g. Post Processing) will need only a few nodes with moderate speed interconnect
- Data movement & storage
  - Data movement between NOAA HPC systems and the Cloud
  - Publicly available data access via internet into the Cloud
  - Majority of data storage in the Cloud and some in NOAA systems
- Workflow
  - All computing components and data movement tied together through workflow

Architecture

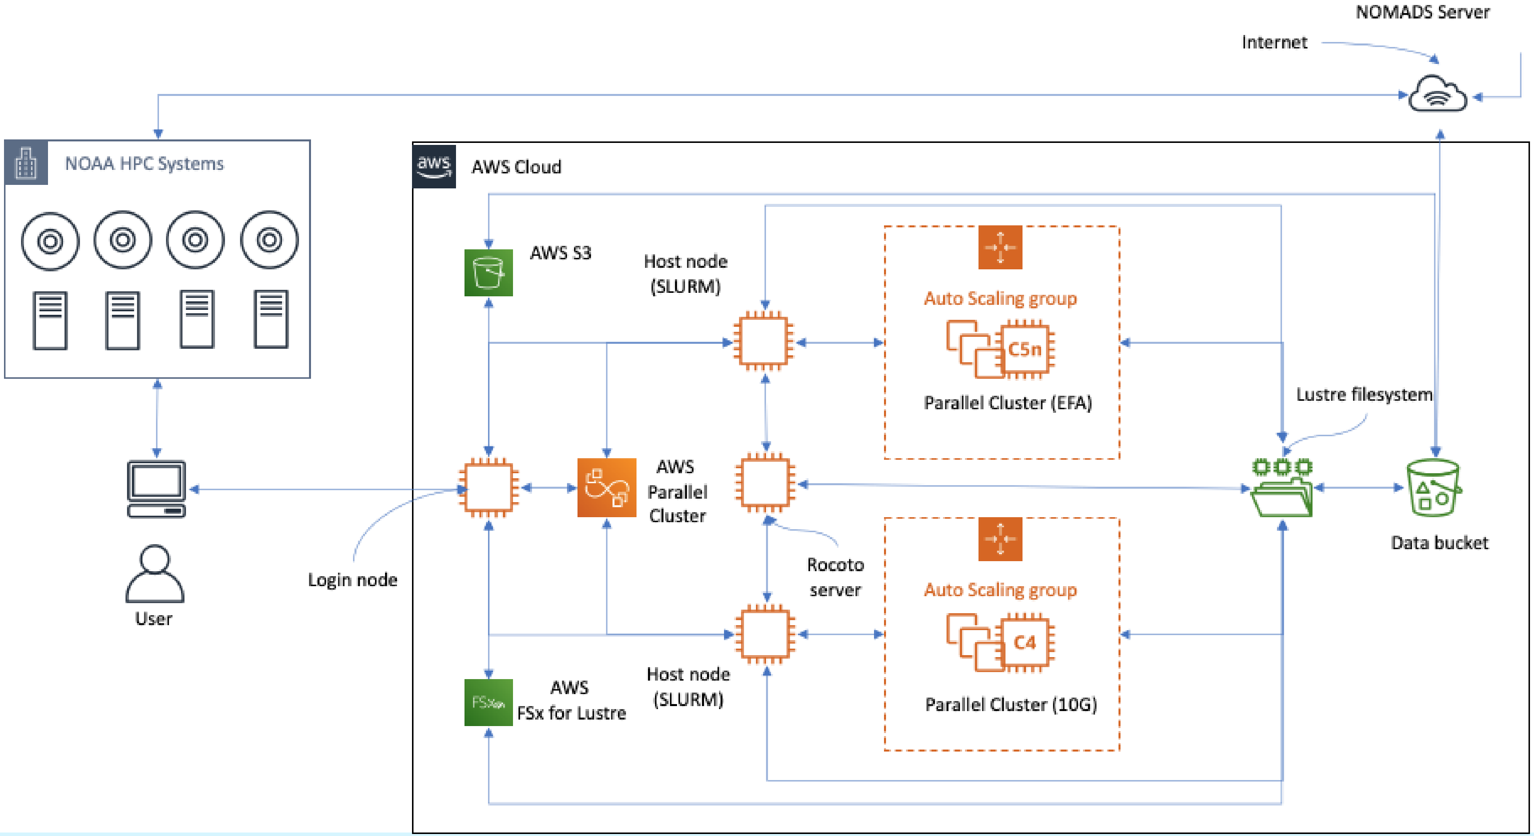

Cloud deployment

- Resource mapping
  - Map each computing component to specific Cloud services/resources – must meet performance or run time requirements
  - Type of instance, choice of file system, storage, etc.
- Cost estimation
  - Cloud offers a wide variety of computing resources with varying cost structures
  - Estimate cost of computing with the chosen resources
- Workflow
  - Ensure enabled workflow meets performance and cost estimates
- Ongoing
  - Cloud Service Providers (e.g. AWS) are constantly revising cost structures, adding new hardware and features several times a year
  - Multi-year project such as RRFS would need periodic assessment & revision of deployment to reduce cost

Resource mapping & cost estimation - UFS example

- Applications built on NOAA systems can run on AWS instances
  - Preliminary work indicated binaries from WCOSS & RDHPCS (Orion, Hera) systems can run in AWS
  - AWS ”c5n” instances and NOAA systems (WCOSS Phase3.5, Orion, Hera) have similar (Intel Xeon Skylake) processors
  - UFS can be run on AWS “c5n” instances with similar performance as on NOAA systems
- Use NOAA systems for performance optimization
  - Much optimization (configuration-related) work can be done on WCOSS & RDHPCS systems – saves spending money on Cloud systems
  - Verification work needs to be done in the Cloud
- Estimate AWS cost for optimized application configuration
  - Cores per node is different: 40 cores/node on WCOSS & RDHPCS while 36 cores/node on AWS ”c5n” instances

UFS - CONUS (details)

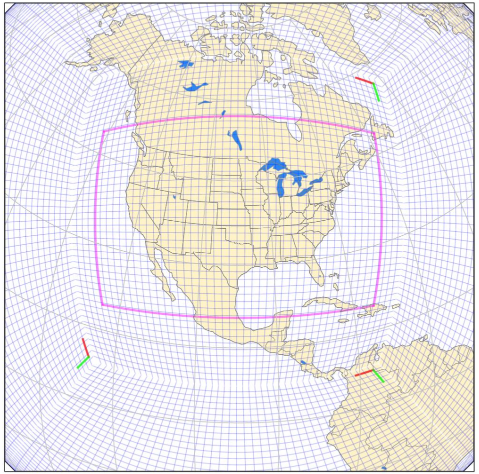

- Grid dimensions (1741 x 1037)
- Time step = 36s
- Forecast duration = 60 hours
- History file output = every hour
- CCPP Physics scheme: FV3_GFS_v15_thomson_mynn

UFS - CONUS (60hr forecast on Hera)

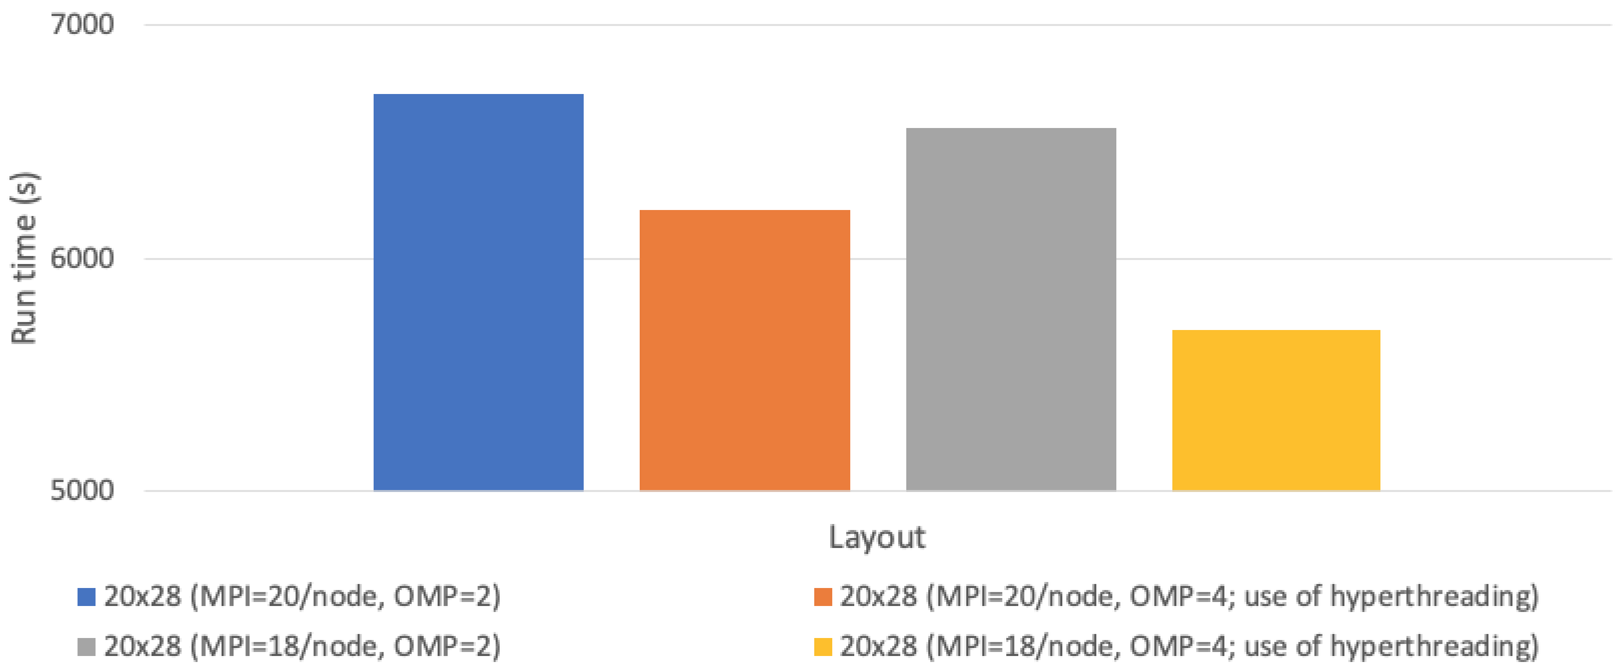

UFS - CONUS 60hr forecast on Hera (node-hours)

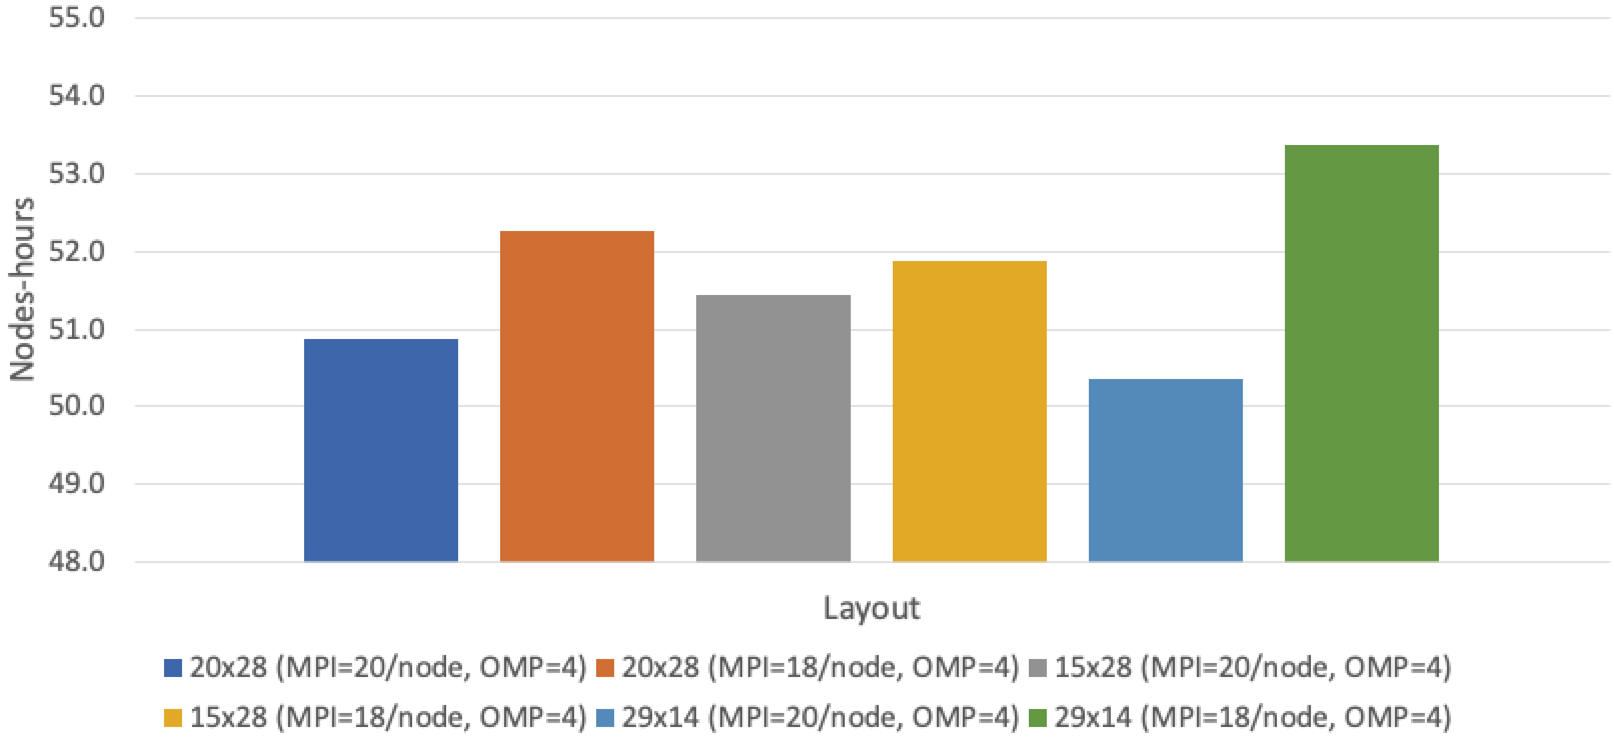

Cost estimation for AWS - unknowns

- Flexibility in run time (e.g. can be > 3hrs) -> use of different instance types
  - Best interconnect (EFA) only comes with more expensive instance types
  - If 10g Ethernet performance is acceptable – may lead to more cost-effective computing
- File system performance
  - Multiple options (parallel Lustre & EFS) available – will require benchmarking
- Processor frequency differences

Prototyping at NOAA Testbeds

Collaborative engagement at NOAA testbeds is a critical part of the development process for RRFS. Some of our cloud resources will be used to execute RRFS prototype components for testbed evaluation.

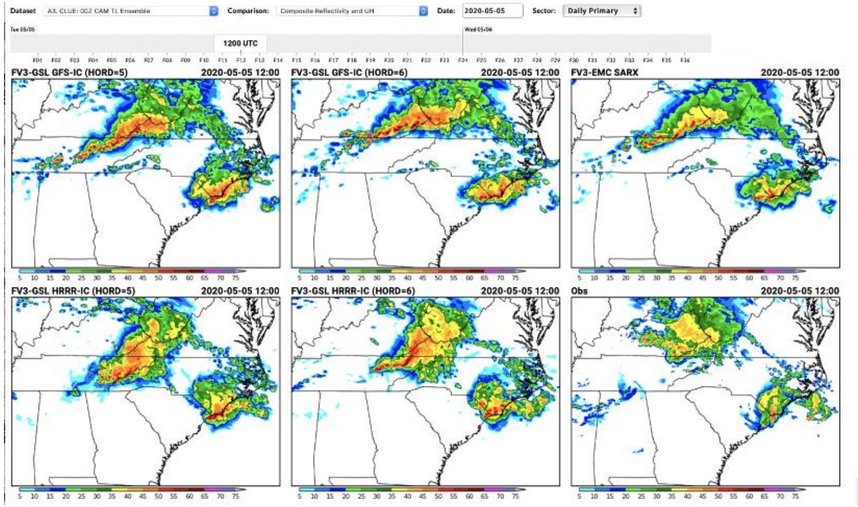

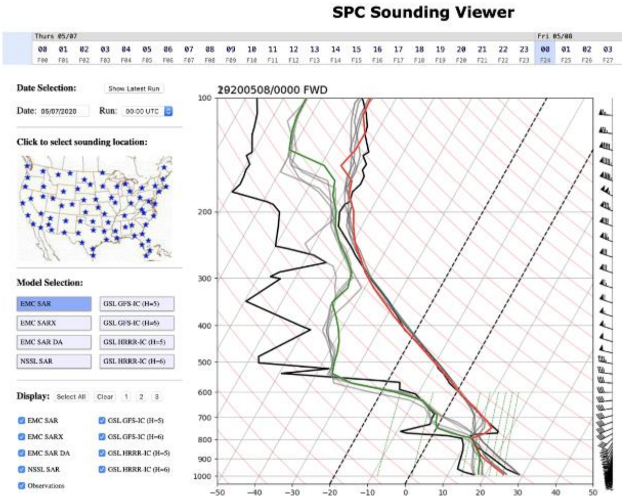

Closing

- RRFSv1 planned for implementation in 2023 timeframe
  - Next generation, rapidly-updated convection-allowing NWP system
- Leveraging the cloud for development
  - Expand beyond conventional compute resources
- Project is in early phases
- Prototype systems to be demonstrated at upcoming NOAA Testbeds (HWT, FFaIR, WWE)